Data management in LCG and EGEE

David Smith

CERN & EGEE-JRA1/SA3 Data Management Team

David.Smith@cern.ch

EGEE Workshop on Management of Rights in Production Grids

www.eu-egee.org

Information Society and Media

Overview

- **Scope of this talk is gLite 3.0 data management components in the context of rights management**
  - gLite 3.0 is a combination software from LCG-2.7, gLite 1.5 and other projects
  - One user community in particular is also using some other components from gLite 1.5
    - Will mention gLite I/O and fireman (file catalog) in connection with BIOMED VO

- **Components in gLite 3.0 relating to rights management:**
  - Encrypted Data Storage (EDS) tools and keystore service
    - Provides encryption and decryption of data
    - Keystore stores the EDS cipher keys
  - Components providing access control list support:
    - EDS keystore
    - The LCG file catalog (LFC)
    - The disc pool manager (DPM)

EDS and the keystore

- **EDS handles encryption/decryption of data**
  - Employs symmetric ciphers via openssl
  - Uses a keystore database via a service called Hydra
  - EDS available as an API. Also as a CLI tool that also manages I/O access via the gLite 1.5 component gLite I/O
    - Will soon provide CLIs for keystore manipulation and encryption/decryption of files without gLite I/O layer
    - In the future will make tools available that have I/O access via GFAL integrated
- **Hydra - The EDS keystore**
  - Is a metadata catalog service
  - Is used to store
    - key, key length, openssl cipher name and cipher IV as necessary
  - Has ACLs on entries to allow fine grained access control
    - Has 3 sets of Perms (8 bits) for user, group, others
    - plus ACLs: Perms on a principal (user DN or VOMS group)

EDS Tools in use

- **Medical community is the main EDS user**
  - Have strict privacy requirements
  - Currently using EDS with glite I/O and fireman
    - gLite I/O and fireman provide a “wrapping” of the storage element (SE) to allow fine grained file access independent of the SE’s functionality
    - Community has their own storage element (SE) called DICOM-SE
  - Files stored on a DICOM-SE are stored in the clear
  - Encrypted before leaving DICOM-SE, so
    - DICOM-SE registers key in Hydra
    - Data are stored on normal SEs on the grid encrypted
    - Decrypted in memory of final application by EDS routines

Example: EDS in use

- **Accessing encrypted data on a standard SE**
  - Note in this example:
  - authorization decision enforcement at the gLite I/O server
    - Ensures fine grained access control to files
  - Encryption also works for output data

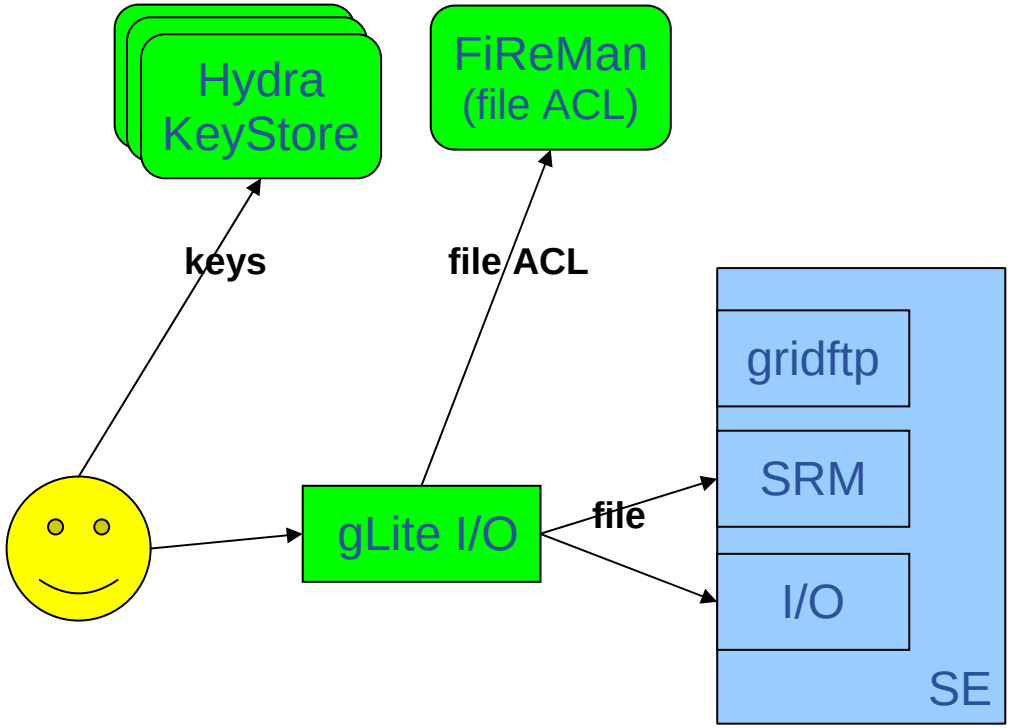

The LHC File Catalog

- **LFC uses ACLs to allow restriction of access to the file catalog**
  - Catalog associates a logical filename (LFN) and unique identifier (GUID) to a file entry
  - File entry holds information on zero or more physical replicas
  - LFNs reside in hierarchical namespace
  - Symbolic links may point LFNs
  - A fixed number of system metadata entries per file
  - A small amount (one field) of user attached metadata per file

Relationships in the Catalog

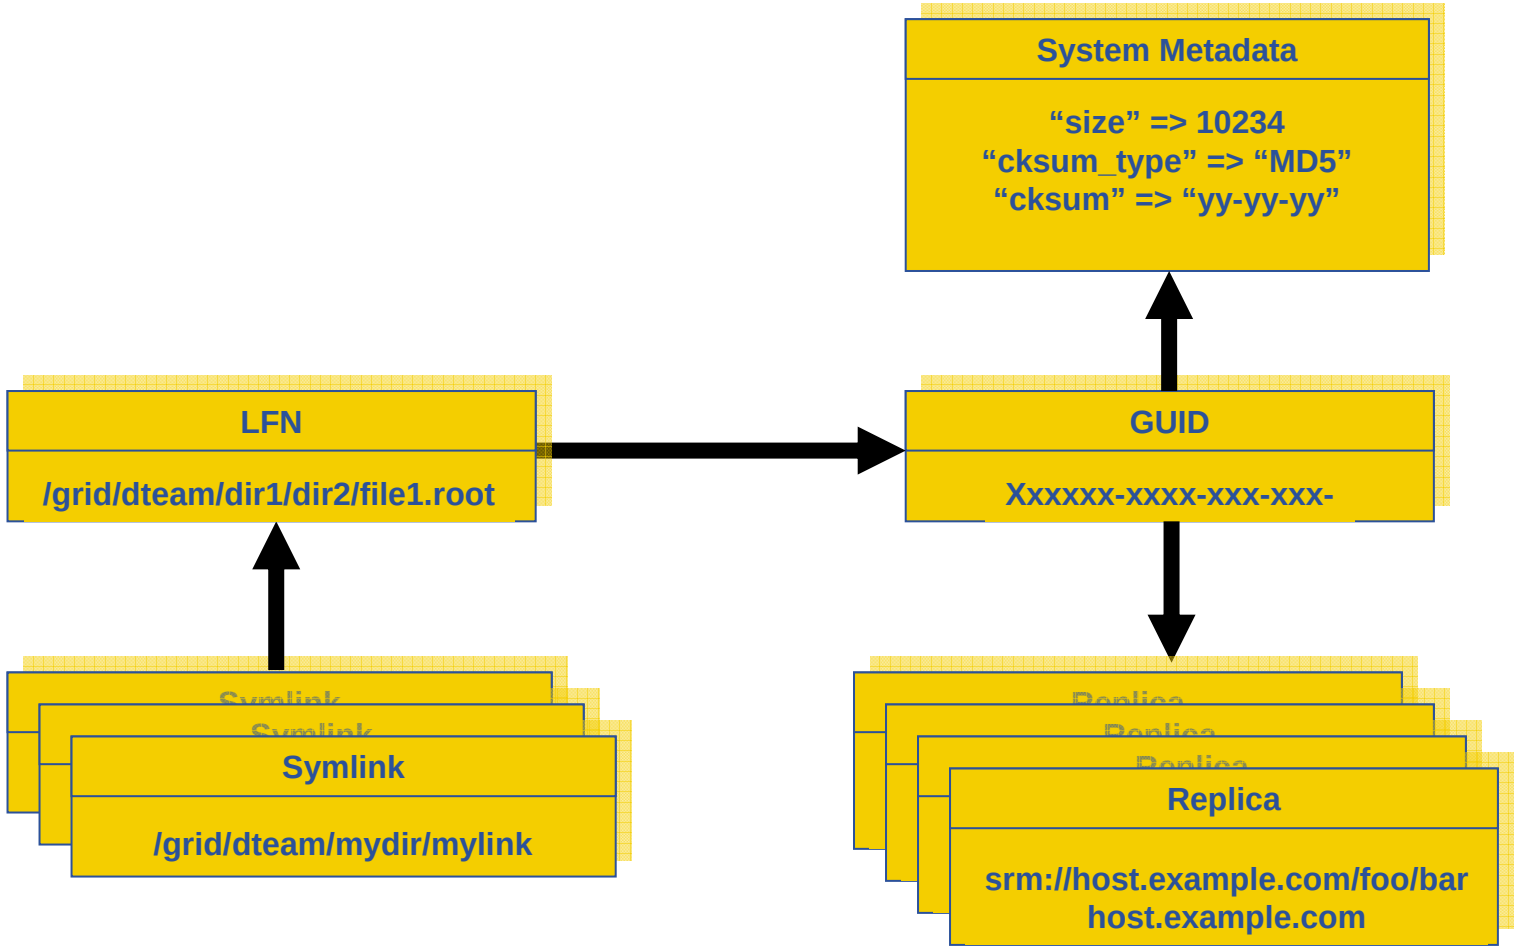

Authorization in the LFC

- DNs are mapped to an internal ID: usually called a virtual ID
- Virtual UID is created on the fly the first time the system receives a request for this DN
- A given user may have one DN and several roles, so a given user may be mapped to one UID and several GIDs
- Currently only the primary role is used
- Support for normal proxies and VOMS proxies
- Administrative tools available to update the DB mapping table:
  - To create VO groups in advance
  - To keep same UID when DN changes
  - To get same UID for a DN and a Kerberos principal

LFC Access Control Lists

- **LFC support Posix ACLs based on virtual ids**
  - Access Control Lists on files and directories
  - Default Access Control Lists on directories: they are inherited by the sub-directories and files under the directory
- **Example**
  - lfc-mkdir /grid/dteam/jpb
  - lfc-setacl -m d:u::7,d:g::7,d:o:5 /grid/dteam/jpb
  - lfc-getacl /grid/dteam/jpb

```
# file: /grid/dteam/jpb
# owner: /C=CH/O=CERN/OU=GRID/CN=Jean-Philippe Baud 7183
# group: dteam
user::rwx
group::r-x              #effective:r-x
other::r-x
default:user::rwx
default:group::rwx
default:other::r-x
```

DPM

- **The LFC forms the *Name Server* component of the DPM**
  - Maps the path found in the SURL to locations within the DPM

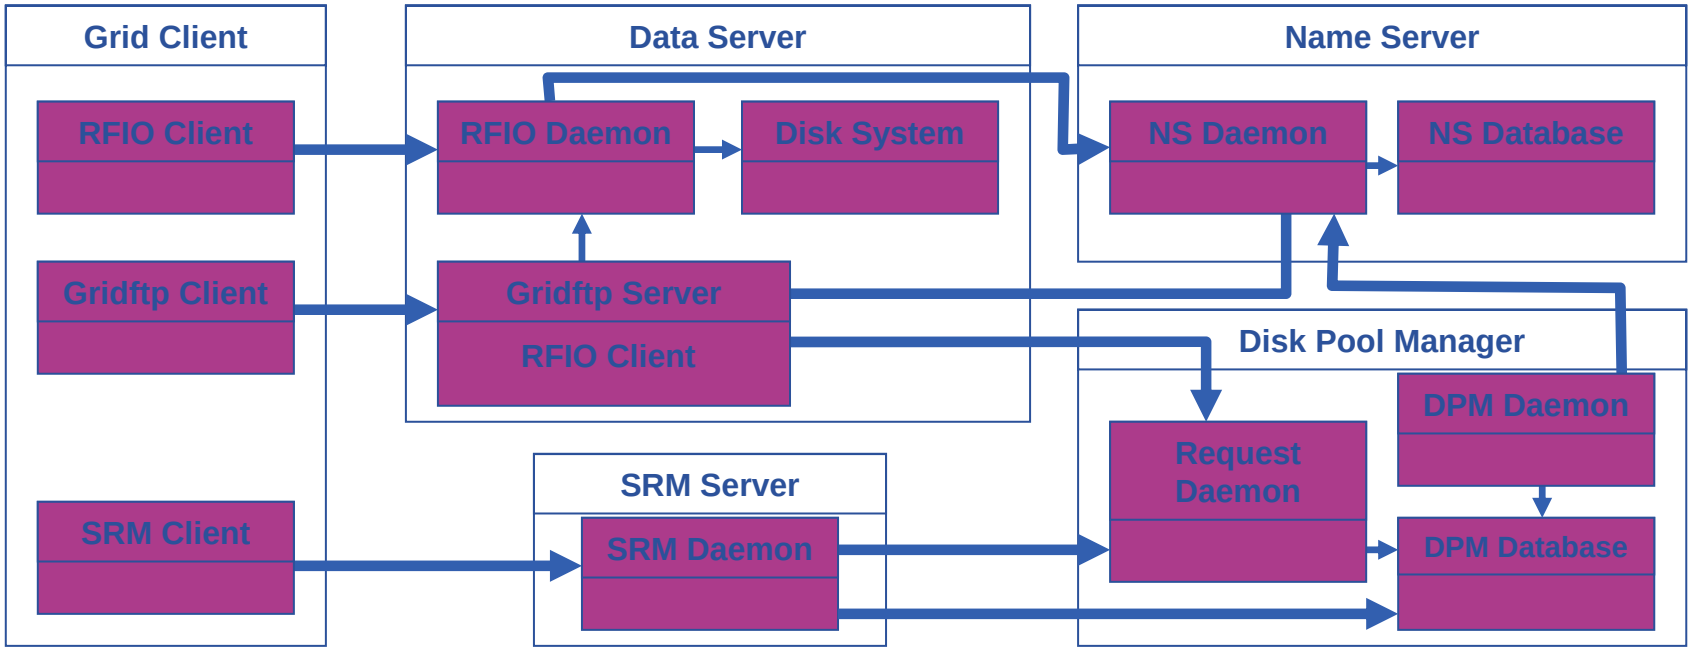

Summary

- **gLite 3.0 and rights management**
  - Encryption of data with EDS
  - ACLs based on user DN and VOMS attributes (if present) in Hydra, LFC and DPM
  - gLite 3.0 does not explicitly provide file level ACLs for arbitrary storage
    - One user community is currently using gLite I/O and the fireman file catalog to do so